

TITLE OF THE INVENTION

**Scaleable Line-Based Protection for Connection Oriented
Communications Protocols**

CROSS REFERENCE TO RELATED APPLICATIONS

--None--

STATEMENT REGARDING FEDERALLY SPONSORED RESEARCH OR
DEVELOPMENT

--Not Applicable--

BACKGROUND OF THE INVENTION

The present invention is related to the field of protection switching in data communications networks.

Many data communications networks employ some form of protection switching to provide better availability of communications services to customers than can be provided by unprotected networks. Generally, protection switching involves the detection of failures within the network, the communication of the failure information to nodes that are affected by a detected failure, and the switching of traffic from one path or connection to another path or connection at the affected nodes as dictated by a predetermined protection switching scheme.

In so-called connection-oriented networks, which employ pre-established virtual and/or physical connections for carrying user data traffic, one class of protection switching schemes is known as "line-based" protection switching. In contrast to source-based schemes, in which traffic is re-routed at its source upon occurrence of a failure and may take a completely different path to its destination, line-based schemes involve more local or hop-by-hop protection switching decisions. Thus, if an end-to-end

connection includes a number of intermediate nodes and connection segments and line-based protection switching is utilized, one or more of the intermediate nodes respond to a failure by taking local actions to re-route the traffic around the failure, without necessarily involving either the source or destination in the protection action. Line-based protection switching can reduce the disruption that can be caused by failures, and under some circumstances may be faster and more efficient than source-based protection switching.

Common examples of both source-based and line-based protection schemes are found in Synchronous Optical Network (SONET) networks. A SONET ring can employ unidirectional path switched ring (UPSR) protection switching or bidirectional line-switched ring (BLSR) protection switching. In UPSR protection switching, information about a failure must be propagated to the destination node, which responds by switching to accept data already flowing on a protect path separate from the working path. However, the destination node may be many hops away from the failure, potentially resulting in a long switchover delay and concomitant loss of data. In BLSR protection switching, the failure information must be propagated to every node on the ring to enable affected traffic to be re-routed in the opposite direction from source to destination, and then each node must perform the necessary switching. In either case, protection switching may be undesirably slow and/or inefficient. Additionally, these techniques suffer relatively poor scalability due to their reliance on relatively wide-area communication of failures and initiation of protection switching actions.

A protection switching technique having improved speed, efficiency and scalability is desirable.

BRIEF SUMMARY OF THE INVENTION

In accordance with the present invention, a scalable protection method for connection oriented networks is disclosed in which protection switching actions are generally taken locally based on pre-established protection connection segments, resulting in improved speed and efficiency in protection switching operations.

The disclosed protected network includes a source and destination, and primary nodes interconnected by working path segments between the source and destination. A number of backup nodes are interconnected by pre-provisioned protection path segments between the source and the destination, and each backup node is also interconnected with an associated primary node by a bidirectional set of pre-provisioned shunt segments.

Each primary node, under normal working circumstances, directs input traffic from an upstream working path segment to a downstream working path segment. Upon occurrence of a failure on the upstream working path segment, a primary node directs input traffic from an input shunt segment to the downstream working path segment, and upon occurrence of a failure on the downstream working path segment, directs input traffic from the upstream working path segment to an output shunt segment.

Each backup node, upon occurrence of a failure on the downstream working path segment of the associated primary node, directs input traffic from an input shunt segment to a downstream protection path segment, and upon occurrence of a failure on the upstream working path segment of the associated primary node, directs input traffic from an upstream protection path segment to an output shunt segment. Additionally if the associated primary node itself fails, the backup node directs input traffic from an upstream protection path segment to a downstream protection path segment.

As a result of these combined operations of the primary and backup nodes when failures occur, new paths are created including shunt segments and protection segments that bypass the failures. Because the shunt and protection segments are pre-provisioned, protection switching is performed rapidly and efficiently. Additionally, the technique is scalable. As primary nodes are added in a network, additional backup nodes and shunt and protection segments can be added in an incremental fashion. Protection switching can be performed in a relatively small neighborhood of a failure, rather than requiring larger-scale communication and switching responses as in present protection architectures.

Other aspects, features, and advantages of the present invention will be apparent from the detailed description that follows.

BRIEF DESCRIPTION OF THE SEVERAL VIEWS OF THE DRAWING

The invention will be more fully understood by reference to the following Detailed Description of the invention in conjunction with the Drawing, of which:

Figure 1 is a block diagram of a network incorporating line based protection in accordance with the present invention;

Figure 2 is a block diagram of switching circuitry in a primary node in the network of Figure 1;

Figure 3 is a block diagram of switching circuitry in a backup node in the network of Figure 1;

Figure 4 is a block diagram of the network of Figure 1 in the presence of a failure on a working network segment;

Figure 5 is a block diagram of the network of Figure 1 in the presence of a failure of a primary node;

Figure 6 is a block diagram of a network incorporating line based protection of a point-to-multipoint connection in accordance with the present invention; and

Figure 7 is a block diagram of the network of Figure 6 in the presence of a failure on a working network segment.

DETAILED DESCRIPTION OF THE INVENTION

Figure 1 shows a portion of an exemplary network for enabling protected unidirectional communication between a source node (S) 10 and a destination node (D) 12. The primary or working communication path includes a number of intermediate nodes designated "primary" nodes (P) 14 interconnected by working segments (WS) 16. Also shown are a set of backup nodes (B) 18, each being associated with a corresponding one of the primary nodes 14. The backup nodes 18 are interconnected between the source node 10 and the destination node 12 by a number of protection segments (PS) 20. Also, each backup node 18 is interconnected with the associated primary node 14 by a corresponding pair of shunt segments (SS) 22. Each pair 22 includes a first shunt segment 21 for carrying traffic from a primary node 14 to the associated backup node 18, and a second shunt segment 23 for carrying traffic from a backup node 18 to the associated primary node 14.

The various segments 16, 20 and 22 are established at the time of connection setup, in advance of carrying any user data traffic from the source node 10 to the destination node 12. The path consisting of the working segments 16 through the primary nodes 14 is a unidirectional path for carrying working traffic from the source node 10 to the destination node 12, and the protection segments 20 are designated to carry protection traffic in the same direction. It will be appreciated that the nodes 10 and 12 may exchange data traffic in the other direction as well (i.e. from node 12 to node 10), for which a separate set of working and protection segments (not shown) must be established. In general, the segments utilized for traffic in the other direction may flow through a different set of primary nodes,

although in practice it is generally advantageous for traffic in both directions to traverse the same set of nodes. Also, a given shunt segment 22 may serve to protect traffic flowing in both directions. In one embodiment, the segments 16, 20, and 22 can be realized as label-switched paths (LSPs) as known in the Multiprotocol Label Switching (MPLS) architecture. They may also be realized as virtual connections (VCs) such as defined in the Asynchronous Transfer Mode (ATM) architecture, or similar pre-established connections.

Figure 2 shows circuitry used for protection switching within the primary nodes 14. First selection circuit 24 selects the source for traffic sent from the primary node 14 on its downstream or output working segment 16, shown as "WS-OUT", and second selection circuit 26 selects the source for traffic sent from the primary node 14 on its output shunt segment 21, shown as "SS-OUT". The inputs to the first selection circuit 24 are (1) the upstream or input working segment 16 ("WS-IN"), (2) the input shunt segment 23 ("SS-IN") and (3) "no connection" 28 (shown as "nc"). When no connection 28 is selected, the output working segment 16 is not being utilized to carry traffic. This case corresponds to the presence of a failure downstream of the primary node 14, as explained below.

The inputs to the second selection circuit 26 are (1) "no connection" 30 and (2) the upstream or input working segment 16 ("WS-IN"). When no connection 30 is selected, the output shunt segment 21 is not being utilized to carry traffic. This case corresponds to the normal working condition, as explained below.

Figure 3 shows circuitry used for protection switching within the backup nodes 18. First selection circuit 32 selects the source for traffic sent from the backup node 18 on its downstream or output protection segment 20, shown as "PS-OUT", and second selection circuit 34 selects the source for traffic sent from the backup node 18 on its output shunt segment 23, shown as

"SS-OUT". The inputs to the first selection circuit 32 are (1) "no connection" 36, (2) the input shunt segment 21 ("SS-IN"), and (3) the input protection segment 20 ("PS-IN"). When no connection 36 is selected, the output protection segment 20 is not being utilized to carry traffic. This case corresponds to the normal working condition, as explained below. The inputs to the second selection circuit 34 are (1) "no connection" 38 and (2) the upstream or input protection segment 20 ("PS-IN"). When no connection 38 is selected, the output shunt segment 23 is not being utilized to carry traffic. This case corresponds to the normal working condition, as explained below.

The circuitry of Figures 2 and 3 is used for protection switching when necessitated by failures within the network. In general, a given backup node 18 and associated protection segments 20 and shunt segments 22 are utilized to route traffic around a failure at or near the primary node 14 with which the given backup node 18 is associated. Specific examples of such failures are given below. From the perspective of a given primary node 14, failures can be categorized as having occurred at the primary node 14 itself, "upstream" of the primary node 14, i.e. toward the source node 10, or "downstream" of the primary node 14, i.e., toward the destination node 12. A failure of a primary node 14 itself is considered to be downstream of a working upstream primary node 14, and upstream of a working downstream primary node 14. This specific scenario is also described below.

Tables 1 and 2 summarize the operation of the switch circuits 24, 26, 32 and 34 at a primary node 14 and associated backup node 18 based on the existence of and relative location of a failure. The contents of these tables are explained below.

Table 1 - Switch Circuits at Primary Node

Case	WS-OUT	SS-OUT
Working	WS-IN	nc
Failure - upstream	SS-IN	nc
Failure - downstream	nc	WS-IN

Table 2 - Switch Circuits at Backup Node

Case	PS-OUT	SS-OUT
Working	nc	nc
Failure - upstream	nc	PS-IN
Failure - downstream	SS-IN	nc
Failure - primary node	PS-IN	nc

10071712.00032
10
20
30
40
50
60
70
80
90
100
110
120
130
140
150
160
170
180
190
200
210
220
230
240
250
260
270
280
290
300
310
320
330
340
350
360
370
380
390
400
410
420
430
440
450
460
470
480
490
500
510
520
530
540
550
560
570
580
590
600
610
620
630
640
650
660
670
680
690
700
710
720
730
740
750
760
770
780
790
800
810
820
830
840
850
860
870
880
890
900
910
920
930
940
950
960
970
980
990
1000
1010
1020
1030
1040
1050
1060
1070
1080
1090
1100
1110
1120
1130
1140
1150
1160
1170
1180
1190
1200
1210
1220
1230
1240
1250
1260
1270
1280
1290
1300
1310
1320
1330
1340
1350
1360
1370
1380
1390
1400
1410
1420
1430
1440
1450
1460
1470
1480
1490
1500
1510
1520
1530
1540
1550
1560
1570
1580
1590
1600
1610
1620
1630
1640
1650
1660
1670
1680
1690
1700
1710
1720
1730
1740
1750
1760
1770
1780
1790
1800
1810
1820
1830
1840
1850
1860
1870
1880
1890
1900
1910
1920
1930
1940
1950
1960
1970
1980
1990
2000
2010
2020
2030
2040
2050
2060
2070
2080
2090
2100
2110
2120
2130
2140
2150
2160
2170
2180
2190
2200
2210
2220
2230
2240
2250
2260
2270
2280
2290
2300
2310
2320
2330
2340
2350
2360
2370
2380
2390
2400
2410
2420
2430
2440
2450
2460
2470
2480
2490
2500
2510
2520
2530
2540
2550
2560
2570
2580
2590
2600
2610
2620
2630
2640
2650
2660
2670
2680
2690
2700
2710
2720
2730
2740
2750
2760
2770
2780
2790
2800
2810
2820
2830
2840
2850
2860
2870
2880
2890
2900
2910
2920
2930
2940
2950
2960
2970
2980
2990
3000
3010
3020
3030
3040
3050
3060
3070
3080
3090
3100
3110
3120
3130
3140
3150
3160
3170
3180
3190
3200
3210
3220
3230
3240
3250
3260
3270
3280
3290
3300
3310
3320
3330
3340
3350
3360
3370
3380
3390
3400
3410
3420
3430
3440
3450
3460
3470
3480
3490
3500
3510
3520
3530
3540
3550
3560
3570
3580
3590
3600
3610
3620
3630
3640
3650
3660
3670
3680
3690
3700
3710
3720
3730
3740
3750
3760
3770
3780
3790
3800
3810
3820
3830
3840
3850
3860
3870
3880
3890
3900
3910
3920
3930
3940
3950
3960
3970
3980
3990
4000
4010
4020
4030
4040
4050
4060
4070
4080
4090
4100
4110
4120
4130
4140
4150
4160
4170
4180
4190
4200
4210
4220
4230
4240
4250
4260
4270
4280
4290
4300
4310
4320
4330
4340
4350
4360
4370
4380
4390
4400
4410
4420
4430
4440
4450
4460
4470
4480
4490
4500
4510
4520
4530
4540
4550
4560
4570
4580
4590
4600
4610
4620
4630
4640
4650
4660
4670
4680
4690
4700
4710
4720
4730
4740
4750
4760
4770
4780
4790
4800
4810
4820
4830
4840
4850
4860
4870
4880
4890
4900
4910
4920
4930
4940
4950
4960
4970
4980
4990
5000
5010
5020
5030
5040
5050
5060
5070
5080
5090
5100
5110
5120
5130
5140
5150
5160
5170
5180
5190
5200
5210
5220
5230
5240
5250
5260
5270
5280
5290
5300
5310
5320
5330
5340
5350
5360
5370
5380
5390
5400
5410
5420
5430
5440
5450
5460
5470
5480
5490
5500
5510
5520
5530
5540
5550
5560
5570
5580
5590
5600
5610
5620
5630
5640
5650
5660
5670
5680
5690
5700
5710
5720
5730
5740
5750
5760
5770
5780
5790
5800
5810
5820
5830
5840
5850
5860
5870
5880
5890
5900
5910
5920
5930
5940
5950
5960
5970
5980
5990
6000
6010
6020
6030
6040
6050
6060
6070
6080
6090
6100
6110
6120
6130
6140
6150
6160
6170
6180
6190
6200
6210
6220
6230
6240
6250
6260
6270
6280
6290
6300
6310
6320
6330
6340
6350
6360
6370
6380
6390
6400
6410
6420
6430
6440
6450
6460
6470
6480
6490
6500
6510
6520
6530
6540
6550
6560
6570
6580
6590
6600
6610
6620
6630
6640
6650
6660
6670
6680
6690
6700
6710
6720
6730
6740
6750
6760
6770
6780
6790
6800
6810
6820
6830
6840
6850
6860
6870
6880
6890
6900
6910
6920
6930
6940
6950
6960
6970
6980
6990
7000
7010
7020
7030
7040
7050
7060
7070
7080
7090
7100
7110
7120
7130
7140
7150
7160
7170
7180
7190
7200
7210
7220
7230
7240
7250
7260
7270
7280
7290
7300
7310
7320
7330
7340
7350
7360
7370
7380
7390
7400
7410
7420
7430
7440
7450
7460
7470
7480
7490
7500
7510
7520
7530
7540
7550
7560
7570
7580
7590
7600
7610
7620
7630
7640
7650
7660
7670
7680
7690
7700
7710
7720
7730
7740
7750
7760
7770
7780
7790
7800
7810
7820
7830
7840
7850
7860
7870
7880
7890
7900
7910
7920
7930
7940
7950
7960
7970
7980
7990
8000
8010
8020
8030
8040
8050
8060
8070
8080
8090
8100
8110
8120
8130
8140
8150
8160
8170
8180
8190
8200
8210
8220
8230
8240
8250
8260
8270
8280
8290
8300
8310
8320
8330
8340
8350
8360
8370
8380
8390
8400
8410
8420
8430
8440
8450
8460
8470
8480
8490
8500
8510
8520
8530
8540
8550
8560
8570
8580
8590
8600
8610
8620
8630
8640
8650
8660
8670
8680
8690
8700
8710
8720
8730
8740
8750
8760
8770
8780
8790
8800
8810
8820
8830
8840
8850
8860
8870
8880
8890
8900
8910
8920
8930
8940
8950
8960
8970
8980
8990
9000
9010
9020
9030
9040
9050
9060
9070
9080
9090
9100
9110
9120
9130
9140
9150
9160
9170
9180
9190
9200
9210
9220
9230
9240
9250
9260
9270
9280
9290
9300
9310
9320
9330
9340
9350
9360
9370
9380
9390
9400
9410
9420
9430
9440
9450
9460
9470
9480
9490
9500
9510
9520
9530
9540
9550
9560
9570
9580
9590
9600
9610
9620
9630
9640
9650
9660
9670
9680
9690
9700
9710
9720
9730
9740
9750
9760
9770
9780
9790
9800
9810
9820
9830
9840
9850
9860
9870
9880
9890
9900
9910
9920
9930
9940
9950
9960
9970
9980
9990
10000

Tables 1 and 2 are explained as follows. At a primary node 14, in the absence of a failure, the traffic from WS-IN is passed along to WS-OUT, and no traffic is sent on SS-OUT, because protection is not active due to the absence of a failure. When a failure occurs upstream of the primary node 14, traffic is still sent on WS-OUT, but the source is the associated backup node 18 via SS-IN. When a failure occurs downstream of the primary node 14, traffic is still received from WS-IN, but is sent to the associated backup node 18 via SS-OUT rather than being forwarded along the working path via WS-OUT. Recall that from the perspective of a given primary node 14, the failure of another primary node 14 is either an upstream or downstream failure, depending on its relative location.

At a backup node 18 (Table 2), in the absence of a failure, no traffic is sent on either PS-OUT or SS-OUT. This is an idle or standby condition. When a failure occurs upstream of the associated primary node 14, the backup node 18 accepts traffic from PS-IN and directs it to the associated primary node 14 via SS-OUT. When a failure occurs downstream of the associated

primary node 14, the backup node 18 accepts traffic from SS-IN and directs it along the protection path via PS-OUT. When the primary node 14 associated with the backup node 18 fails, then traffic is accepted from PS-IN and directed along the protection path via PS-OUT.

Figure 4 depicts the operation of the network in the presence of a failure on the working segment 16-2 extending between two primary nodes 14-1 and 14-2. At the primary node 14-1, the traffic is directed from the input working segment 16-1 toward the backup node 18-1 along the shunt segment 21-1. The backup node 18-1 accepts the traffic from the shunt segment 21-1 and directs the traffic toward the downstream backup node 18-2 along the protection segment 20-2. From the perspective of the backup node 18-2 and the primary node 14-2, the failure is an "upstream" failure. Therefore, the backup node 18-2 directs traffic from the protection segment 20-2 toward the primary node 14-2 via the shunt segment 23-2, and the primary node 14-2 accepts the traffic from the shunt segment 23-2 and directs it to primary node 14-3 via the working segment 16-3.

Figure 5 shows operation when a primary node such as primary node 14-2 fails. In this case, operation of nodes 14-1 and 18-1 is the same as for the situation of Figure 4, and nodes 14-3 and 18-3 operate in the same fashion as do nodes 14-2 and 18-2 in the situation of Figure 4. Additionally, backup node 18-2 forwards traffic from its input protection segment 20-2 toward the downstream backup node 18-3 via output protection segment 20-3. As a result, traffic is routed around failed primary node 14-2.

The preceding description has focused on point-to-point connections having one source node 10 and one destination node 12. The disclosed protection technique can also be utilized in connection with point-to-multipoint connections having a single source and multiple destinations.

10061712.020702

5 Figure 6 shows an example of a point-to-multipoint connection on which the source node 10 sends data to two different destinations 12A and 12B. In this simple two-destination connection, the primary node 14-2 is responsible for replicating the traffic on two output working segments 16-3A and 16-3B, and likewise the backup node 18-2 is responsible for replicating the traffic on two output protection segments 20-3A and 20-3B. The nodes 14-2 and 18-2 are referred to herein as a "branching primary node" and "branching backup node" respectively. The nodes 14-2 and 18-2 operate as shown in Figures 2 and 3 with respect to both the traffic stream for destination 16A and the traffic stream for destination 16B. Generally, it is preferred that the protection switching for these different streams be carried out independently, so that for example a failure of primary node 14-3B would result in protection switching occurring for the traffic for destination 12B but no protection switching occurring for the traffic for destination 12A.

25 Figure 7 shows the existence of a failure on the working segment 16-3A of the "A" branch of the point-to-multipoint connection. In this case, the traffic destined for destination node 12A is directed along shunt segment 21-2 to backup node 18-2, then along protection segment 20-3A to backup node 18-3A, and then along shunt segment 23-3A to primary node 14-3A, which forwards the traffic to destination node 12A along working segment 14-4. The traffic destined for destination node 12B is not affected by this failure, and continues to flow along working segments 16-3B and 16-4B.

30 While in the illustrated embodiments, there is a different backup node 18 associated with each primary node 14, in alternative embodiments a node may serve as a backup node 18 for two or more primary nodes 14, as long as the necessary working segments, protection segments, and shunt segments can be established. It is generally preferred for reliability reasons

that a primary node be directly connected to its associated backup node, although it is not strictly required. By "directly connected", it is meant that there are no intervening nodes that terminate network segments such as LSPs. A lower-level device such as an electrical repeater or hub would generally not qualify as an intervening node. As already mentioned, there may be additional nodes within one or more of the protection segments 20 that do not participate in the protection operation as a backup node 18. Additionally, it is possible that such additional nodes are also included within the working segments 16, although such configurations are preferably avoided. Generally, it is preferred that each node along the working path from source 10 to destination 12 be protected.

It will be apparent to those skilled in the art that modifications to and variations of the disclosed methods and apparatus are possible without departing from the inventive concepts disclosed herein, and therefore the invention should not be viewed as limited except to the full scope and spirit of the appended claims.